# DS-DF interface document

Giles Barr – thoughts from the data selection side

The document is on EDMS at

https://edms.cern.ch/ui/#!master/navigator/project?P:100305685:100652014:subDocs

We have put a copy on DUNE docdb-21946

Text in red (and final slide) is notes added during talk

# Introduction

- There are draft interface documents on EDMS.

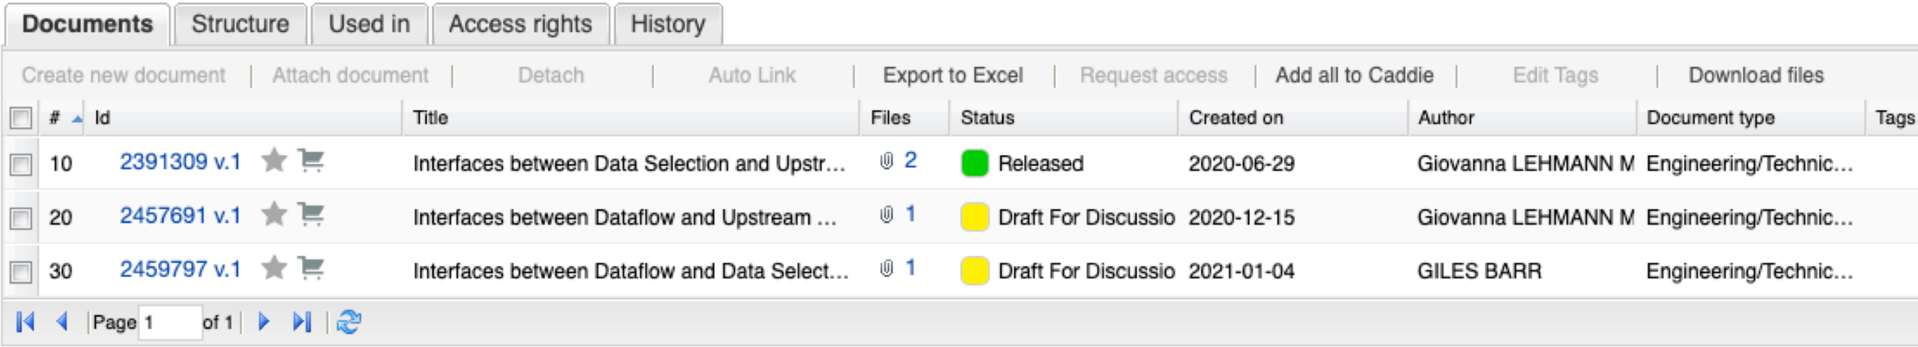

| □ | # | Id | Title | Files | Status | Created on | Author | Document type | Tags |
|---|---|---|---|---|---|---|---|---|---|
| □ | 10 | 2391309 v.1 | Interfaces between Data Selection and Upstr… | 2 | 🟩 Released | 2020-06-29 | Giovanna LEHMANN M | Engineering/Technic… | |
| □ | 20 | 2457691 v.1 | Interfaces between Dataflow and Upstream … | 1 | 🟨 Draft For Discussio | 2020-12-15 | Giovanna LEHMANN M | Engineering/Technic… | |
| □ | 30 | 2459797 v.1 | Interfaces between Dataflow and Data Select… | 1 | 🟨 Draft For Discussio | 2021-01-04 | GILES BARR | Engineering/Technic… | |

- The UD-DF one is by far the best (i.e. the one we are not involved in!) – but that is because there are multiple hand-overs being done there.

# DS-DF

- Connection0:
  - The data from upstream DAQ (trigger primitives from TPC, trigger primitives from photodetectors, ‘trigger primitives’ from external trigger sources, such as calibration). All arrive on AppFwk-standard streams/IPC.
- Connection1:
  - Reporting selected triggers to dataflow
- Connection 2:
  - Providing diagnostic/traceback information about each trigger to datafile for offline analysis (trigger efficiency etc)
- Connection 3:
  - High-level filter interface

# Connection 0:

- Diagram below is from UD-DF interface document. It shows the various flows and stores of data in UD,
- Each has a connection with dataflow.
- Data selection on the other hand has one connection with UD and it goes via a standard interface for which dataflow are responsible (appfwk).

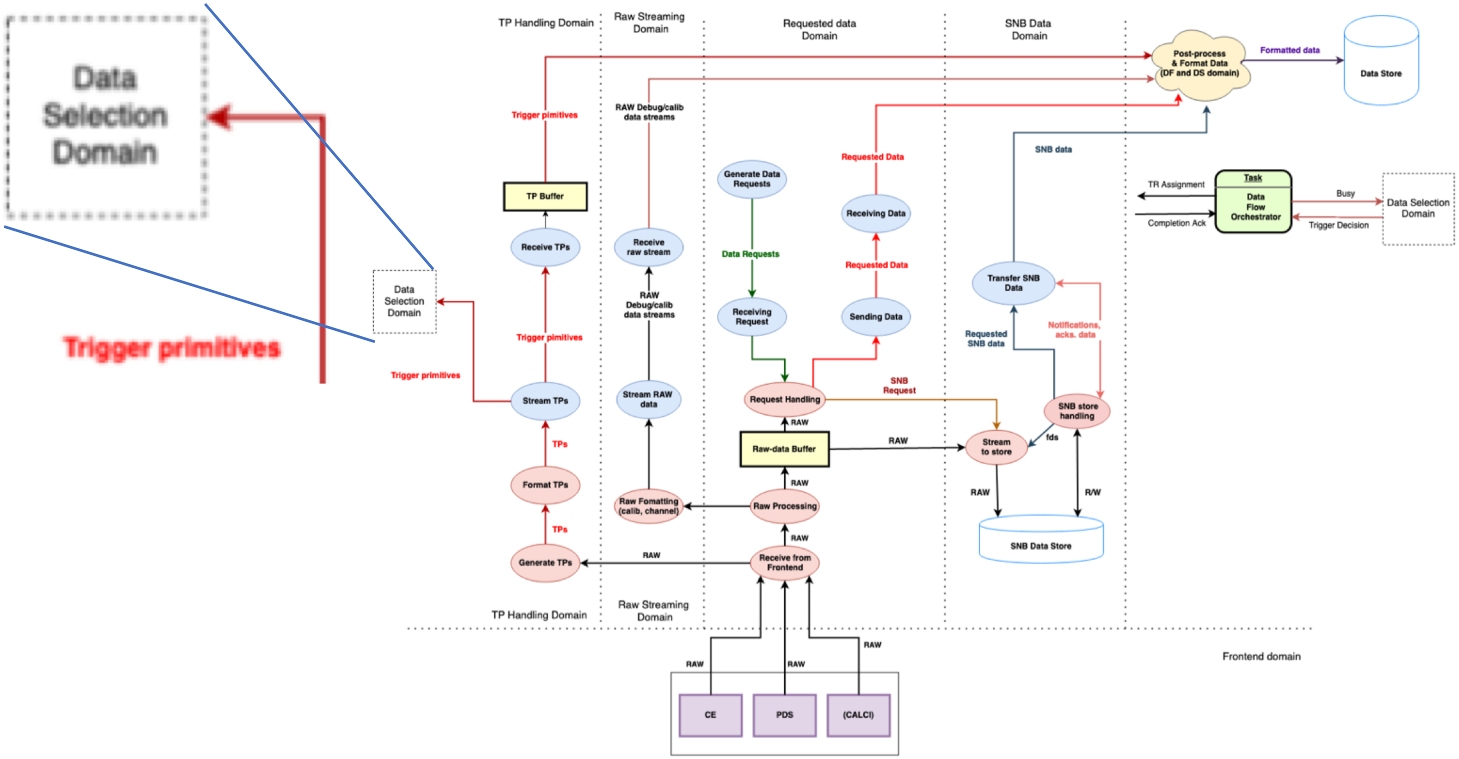

- Connection 1:
- Each trigger shall be communicated as a separate message.
- The trigger is distributed with a DF-streamer structure containing information about each trigger (next slide). The hand-over of messages includes a busy signal sent back from DF to DS if the downstream part of the system will soon block. (Throttling mechanism for when there is a configuration mistake etc.)
- Triggers do not need to be time ordered, they will be sent to dataflow as soon as we are 100% sure they are triggers.
- The data selection will be allowed to send triggers that duplicate readout of data blocks from partially-overlapping events. (e.g. calibration occurs near solar-neutrino event). Data-selection should avoid 100% overlapping triggers, dataflow can accept them.
- There can only be at most one active SNB trigger at a time

The sentence in red generated a lot of discussion. See notes added on last slide

- Connection 1: (cont)
- Trigger info:
- Trigger selection bits (a.k.a.’type’) dataflow only cares if it is SNB or not.   HLF will care in a more granular way, e.g. ‘is it calib’
- Readout mask of APAs (geographical), or groups of PDS (and sub-detectors generally)
- Trigger ID (not guaranteed to increment sequentially)
- Desired readout window and trigger time (three numbers). (a) Time trigger occurred, (b) start of window (c) end of window

- Connection 2: Tracking & audit information requests

- Dataflow intends to provide a ‘streamer’ interface. Collect continuous stream of timestamped data. When triggers arrive, pack data for event building.
  - all these functions within dataflow part of system.

- This DF-streamer is extremely useful for us. Will homogenize this task across the DAQ (in the same way the board-reader interface does in artDAQ).

- We may need to offer volunteers to help.

## Connection 3: High level filter interface

- HLF runs algorithms provided by data selection WG.
- Dataflow WG will deliver trigger records to the HLF processes and will collect results of filtering operation and formatting of data for delivery offline.
- Dataflow components will provide information about available trigger records to HLF and the HLF process can choose to access full or partial events. Data flow will provide tools to access the data. The anticipated granularity of partial events is at the FELIX link level. [Also allows access to Tracking/audit info from earlier part of trigger, and the trigger information from connection1:]
- HLF processes will be able to filter out complete trigger records, accept complete trigger records, accept a particular trigger record but specify that only a part of the record be sent offline,. HLF processes will also be able to add data such as algorithm results to the trigger records.
  - Need to avoid proliferation of formats
  - Reformatting will not be done on the important ‘day-one’ triggers
  - It is intended for a solar neutrino analysis where we can take many small events. Maybe the requirement is that whoever writes that part of the HLF also writes the offline reception routines for these data blocks.

There was a long discussion about whether it is possible for the connection1 output (the MLT sending information to the DFO) to be out of order. The assumption in the earlier designs is that everything should be time ordered throughout the system. However, it is understood that the dataflow goes into parallel processing immediately after the DFO, so the upstream DAQ does not receive requests in order (because the event building is not running in one place). The advantage for allowing out of order triggers is to allow CPU/GPU intensive triggers (such as deep learning) to take longer than the simpler triggers). Here is the discussion.

- 16:56:34 From Kurt Biery to Everyone : I will confirm with offline that it is OK to have events that are not in time order (or trigger record number order) in the output files)
- 16:57:44 From Tom Junk to Everyone : I'm pretty sure offline doesn't care. But the data within a SNB trigger record ought to be ordered.
- 16:58:39 From Kurt Biery to Everyone : Tom, yes, absolutely! Time slices should be in time order.
- 16:59:26 From Georgia Karagiorgi to Everyone : Along the same lines, Kurt could you please also ask about whether they care if trigger decisions which appear too close to each other (within 5.4ms) should be merged into a single trigger record or not? my assumption would be not but perhaps this is not an assumption shared by everyone
- 17:01:01 From Tom Junk to Everyone : I am less sure about the right answer here. The offline is set up to process trigger records of fixed length, but that sounds like an artificial deficiency that should be possible to code around.
- 17:01:41 From Georgia Karagiorgi to Everyone : All the experiments I know of assume a fixed trigger time t0 per trigger record
- 17:01:57 From Kurt Biery to Everyone : Yes, I will ask about overlaps, too.
- 17:02:08 From Georgia Karagiorgi to Everyone : the length is variable, but only one candidate (triggered) interaction is present. Of course, this is not the case for SNB, by construction
- 17:02:14 From Tom Junk to Everyone : If trigger records are longer and contain overlaps, one can always slice them offline.
- 17:02:17 From Georgia Karagiorgi to Everyone : So perhaps DUNE offline will be broadly flexible that way
- 17:02:35 From Tom Junk to Everyone : Hopefully. We did that slicing for 35t but it was awkward.
- 17:02:35 From Georgia Karagiorgi to Everyone : I personally think we should keep things simple otherwise our HLF will have to slice things also
- 17:03:24 From Tom Junk to Everyone : You get more chances at it offline, so simplifying HLF sounds like a good idea.